

Automated Landmark Identification for Spacecraft Navigation

Robert W. Gaskell
MS 301/150
Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Drive
Pasadena, CA 91109
(818) 354-2116

ABSTRACT

An integrated approach to surface relative optical landmark tracking for spacecraft is being developed. Landmarks are defined as full digital topography/albedo maps and are determined from previous imaging and navigation data. Initially, this technique will speed up ground based optical navigation. Ultimately, it will enable on-board trajectory determination during orbital and landing maneuvers.

1. Introduction

With the prospect of long term orbital as well as landing missions to small bodies, the use of landmarks as both control points and navigation aids is becoming more important. Not only are landmarks invaluable for the determination of the body's shape, but they provide navigation tie points that are fixed to the body itself.

The NEAR mission was the first to exploit landmarks for navigation. Overlays outlining predicted crater rims were aligned with the actual images, the residuals providing an additional input to the orbit determination software. The effort was entirely ground-based and very labor intensive, with over 34,000 images analyzed by hand [1].

In recent years, a great deal of progress has been made in the digital identification and alignment of landmarks. Early efforts in this area [2] used landmarks which were essentially "painted" onto the surface. They had no topographic component.

Now, landmarks are three dimensional objects that snap into place on the surface. In figure 1, three such landmarks, 100x100 pixel maps at 25 meter resolution, are exhibited on a shape model of Phobos. The shape model itself was constructed out of very large maps, such as the 100 meter resolution map shown in figure 2.




Figure 1. Landmark patches on Phobos shape model




Figure 2. Large map used in shape determination.

The software which generated these maps from the imaging data is quite complex. It uses landmarks to register (align) the images, and possibly provide inputs for orbit determination, while at the same time modeling the landmark's albedo and topography. Once a map is created, however, its alignment with an image is easily accomplished.

The separation of tasks into landmark construction and landmark alignment suggests a scenario for autonomous navigation in cases where the surface can be characterized over a long period of time. These include missions such as NEAR, which provided a long orbital phase during which observations were made, and return missions such as those to Mars, where imaging data is already available.

Most of the processing will be performed on the ground. This includes the determination of camera pointing and spacecraft trajectory, discussed in section 2, and the creation of maps describing the landmarks and any necessary shape models, discussed in section 3. Of necessity, this is an iterative process. Better pointing and navigation information leads to better map solutions. Better map solutions lead to better landmark alignment. Better

landmark alignment leads to better pointing and navigation knowledge.

The proposed on-board processing is discussed in section 4. It will be confined to a small portion of the already existing software which correlates illuminated maps with imaging data projected into the same coordinate system. Section 5 will include a discussion of possible applications of this technology, including comet and asteroid missions, and approach and orbital navigation at Mars.

2. Navigation and pointing

A landmark is defined as a three dimensional patch of the target body's surface which contains distinctive features, due to albedo and/or topographic variations. It is specified by a vector \mathbf{V} from the center of a body fixed coordinate system to the origin of a local surface (map) coordinate system, by the axes of that coordinate system \mathbf{u}_i ($i=1,3$), roughly in the south, east and vertical directions respectively, and by heights and albedos at positions (x,y) relative to that system. Thus, coordinates of a body-fixed point \mathbf{P} on the surface are $x=\mathbf{u}_1 \bullet (\mathbf{P}-\mathbf{V})$ and $y=\mathbf{u}_2 \bullet (\mathbf{P}-\mathbf{V})$ with the height $h(x,y)=\mathbf{u}_3 \bullet (\mathbf{P}-\mathbf{V})$ and an albedo $a(x,y)$.

The spacecraft position vector \mathbf{W} and the camera coordinate system \mathbf{c}_i ($i=1,3$), unit vectors in the sample, line and boresight directions, respectively, are also specified in body-fixed coordinates.

Leaving aside the transformation from inertial space to body-fixed coordinates, the potential unknowns in this system are \mathbf{V} and h , i.e. the shape/topography problem, and \mathbf{W} and \mathbf{c}_i , i.e. the navigation/pointing problem.

These two problems are related by the images taken from the spacecraft. For a simple narrow angle camera with focal length f , the image location X,Y on the focal plane of the point $\mathbf{P}(x,y,h)$ on the surface is given by

$$1a) \quad X = f(\mathbf{P}-\mathbf{W}) \bullet \mathbf{c}_1 / (\mathbf{P}-\mathbf{W}) \bullet \mathbf{c}_3$$

$$1b) \quad Y = f(\mathbf{P}-\mathbf{W}) \bullet \mathbf{c}_2 / (\mathbf{P}-\mathbf{W}) \bullet \mathbf{c}_3$$

which can be expanded to

$$2a) \quad X = f[(\mathbf{V}-\mathbf{W}) \bullet \mathbf{c}_1 + M_{11}x + M_{12}y + M_{13}h] / [(\mathbf{V}-\mathbf{W}) \bullet \mathbf{c}_3 + M_{31}x + M_{32}y + M_{33}h]$$

$$2b) \quad Y = f[(\mathbf{V}-\mathbf{W}) \bullet \mathbf{c}_2 + M_{21}x + M_{22}y + M_{23}h] / [(\mathbf{V}-\mathbf{W}) \bullet \mathbf{c}_3 + M_{31}x + M_{32}y + M_{33}h]$$

where $M_{ij} = \mathbf{c}_i \bullet \mathbf{u}_j$. The vector \mathbf{V} to the origin of the landmark map specifies the body-fixed landmark location. Its location in an image is found from equations 2, with x , y , and h set equal to zero. The terms involving M_{ij} play a role in the construction and alignment of landmark maps, and will be discussed in the next section.

If the ensemble of landmark vectors were precisely known, as well as their locations (X,Y) in an image, then the camera pointing \mathbf{c}_i , and the spacecraft vector \mathbf{W} could be determined. An error in pointing is characterized by a translation or rotation of the landmarks in the image. The landmarks maintain a rigid pattern, their angular separations remaining constant. An error in the spacecraft vector is reflected in a change in the angular separation of the landmarks. A change in range is indicated by an overall increase or decrease of angular separation, while translation in another direction produces differential shifts in that direction. The solution for spacecraft position degrades with distance, since it depends upon the separation of landmark points relative to the range.

Similarly, if the camera pointings and spacecraft locations were precisely known for the ensemble of images, as well as the locations (X,Y) of a landmark in those images, then the body-fixed vector \mathbf{V} to that landmark could be determined.

It is almost, but not quite, possible to obtain a complete solution for spacecraft location, camera pointing and landmark location from the imaging data alone. Some other data is required, and this is envisioned as only part of the complete navigation package. First of all, the optical data alone needs to have a scale defined. Clearly, the spacecraft the spacecraft could be twice as far away, and moving twice as fast, and the body could be twice as large. Only radio data can set the scale. Second, notice that only the difference between \mathbf{V} and \mathbf{W} occurs in equations 2. For this reason, the origin of the body-fixed coordinate system is not necessarily at the center of figure. That determination must wait until the landmarks are tied to a shape model. Whether the center of figure is the same as the center of mass must again be determined by examining the doppler data. Finally, the question of transformation between inertial and body-fixed coordinates must ultimately depend upon the precise location of landmarks. This simply introduces another unknown into the iteration loop.

3. Topography and shape

The vector \mathbf{V} is clearly somewhat arbitrary. A vector $b\mathbf{u}_3$ can be added to it if, at the same time, b is subtracted from all of the heights $h(x,y)$. So the problem is really the determination of the heights, given an assumed vector \mathbf{V} to the origin of the map. This is accomplished using stereo-photoclinometry, essentially using all the imaging data available to reconstruct the scene.

For fixed \mathbf{W} and \mathbf{c}_i , and with a specified \mathbf{V} and \mathbf{u}_i , the (X,Y) location of a surface point $\mathbf{P}(x,y,h)$ in map coordinates can be determined for each image. The corresponding brightnesses, drawn from the original images, can be displayed as a function of x and y . Such a display, corresponding to the bottommost landmark in figure 1, is shown in figure 3.




Figure 3. Image data extracted and projected into map coordinates.




Figure 4. Illuminated map from height, slope and albedo data.

If the landmarks are already positioned close enough, to within a few percent of the size of the map, then the brightness can be fit to a simple function of slope and albedo

$$3) \quad I(x,y) = I_0(k) (1+t_3(x,y)) F(\cos i, \cos e, g)$$

where F is an appropriate reflectance function [3] .

During its projection to map coordinates, the imaging data is maximally stretched. The multiplier $I_0(k)$ is an overall scaling factor to align the stretched data from image k with the model to be constructed. The factor $(1+t_3)$ is the relative albedo at map coordinates (x,y) , g is the phase angle, i is the local angle of incidence and e is the local angle of emission. In terms of the slopes $t_1 = -dh/dx$ and $t_2 = -dh/dy$, the local arguments of the function F are

$$4a) \quad \cos i = (s_1 t_1 + s_2 t_2 + s_3) / \sqrt{(1+t_1^2+t_2^2)}$$

$$4b) \quad \cos e = (e_1 t_1 + e_2 t_2 + e_3) / \sqrt{(1+t_1^2+t_2^2)}$$

where \mathbf{s} and \mathbf{e} are the map system components of the sun and spacecraft unit vectors respectively.

Notice that if the alignment is not good enough for this procedure to converge, then the scale can be increased until it is. Once a rough alignment is achieved, the scale can be reduced to the desired value.

Each map is typically about a hundred pixels across. At each of these pixels, the values of t_i are solved for, fitting the brightness data from the images in a least squares sense. After determining the global factor I_0 for each picture, reconstructed images are made using equations 3) and 4). A reconstruction of the images of figure 3 is shown in figure 4.

If the camera pointing \mathbf{c}_i , spacecraft position \mathbf{W} , or landmark location \mathbf{V} is not correct, the actual and predicted images will not be aligned. Correlation of the map images is used to determine the shift and this, in turn, is translated into revised (X,Y) landmark locations in the actual images. After an ensemble of landmark locations has been determined in this manner, equations 2 are solved iteratively for \mathbf{c}_i , \mathbf{W} and \mathbf{V} .

The landmark vector \mathbf{V} is only roughly determined at this stage because the heights are still unknown. They may have been set to zero initially, or they may have been estimated from the current best shape model. In either case, the original images of figure 3 will be distorted relative to the maps of figure 4 since

if a height is in error, equations 2 will extract data from the wrong (X,Y) of the original images.

The distortion due to incorrect heights can be used to correct those heights at some points in the map. A small patch around (x,y) in the reconstruction is compared with the same patch of extracted data as the assumed height is varied. If a sufficiently strong peak is observed, then $h(x,y)$ is taken to be the corresponding height. The number of heights determined in this way depends upon the complexity of the scene, and upon the original imaging geometry.

Starting with this sparse set of stereo points, the array of heights is now completed by integrating the slopes determined in the fit to Equation 3. For example, if the height is known at the point (i-1,j) on the pixelized map, then the height at (i,j) is

$$5a) \quad h(i,j) = h(i-1,j) - (t_1(i-1,j) + t_1(i,j))s/2,$$

where s is the pixel spacing. Similarly, if the height is known at (i,j-1),

$$5b) \quad h(i,j) = h(i,j-1) - (t_2(i,j-1) + t_2(i,j))s/2.$$

If more distant heights are known, then these too can give values for $h(i,j)$, and several different paths can be used in the integrations. In practice, more than four hundred possible paths leading from neighboring heights are averaged to obtain $h(i,j)$.

The process above is repeated several times to yield a set of heights consistent with both the stereo and slope data. By minimizing the rms difference between the original stereo heights and the height solutions at those points, any overall multiplicative bias in the slopes, due to a misdetermination of the $I_0(k)$, is corrected. If other data, such as heights from neighboring maps or from MOLA observations are available, they can also be used to constrain the fit.

4. On-board image correlation

Once a map has been constructed using the procedure in the last section, it can be correlated against data in other images to help determine camera pointing and spacecraft location. It should be stressed that the necessary alignment mentioned above, i.e. a few percent of the size of the map, relates to the construction of the map and not its correlation with data from another image.

A 100x100 pixel map, for example, would be characterized by a minimal header containing \mathbf{V} , \mathbf{u}_I , minimum and maximum slopes, albedos and heights, and any necessary labels. To this would be added byte data representing t_1 , t_2 , t_3 , and h . Altogether, the file would be about 320 kilobits. At a typical DS1 data rate of about 2000 bits per second, such a map could be uploaded in less than three minutes, even without compression.

Once on board, the map would be illuminated using equations 3) and 4). Then, using the nominal values of \mathbf{W} and \mathbf{c}_i , along with \mathbf{V} , \mathbf{u}_I , and the heights from the uploaded file, data would be extracted from the image according to equations 2). The correlation of the two images, either by brute force if they are expected to be closely aligned, or otherwise by FFT, would yield an offset which would be translated into a sample/line offset in the original image, the usual observable in the optical navigation software.

The set of maps should be dense enough so that at least three are visible, on average, in an image. This would provide enough information to solve for both camera pointing and spacecraft location. It would also be desirable if the camera were slewed to image landmarks at different distances in the same field of view, providing cleaner separation between pointing and trajectory solutions.

5. Discussion

The use of many pictures in the construction of a map leads to a super-pixelized image which can be easily correlated with images at five times the original resolution. Thus, during the orbital phase of a small body landing mission, maps can be created which can be used during the final approach. It is hoped that this can be tested next year, using NEAR data for simulated approach and landing navigation.

Since this procedure does not require cratered surfaces, but only that the surface be interesting, it should be useful during approach and landing on comets as well. It is important to stress that landmark navigation is relative to the body-fixed system, the same system in which potential landing sites are identified.

Typically, the rms residuals in a landmark correlation are less than 0.2 map pixels. For a 1000x1000 pixel image, this implies that the camera's twist can be determined to better than

.2 mrad, assuming perhaps five landmarks in the field of view. If the pixel size is .05 mrad, corresponding to a 3° field of view, then the error in clock and cone angle solutions should be less than .01 mrad.

A network of navigation maps for Mars is currently being constructed, using Viking, MOC and MOLA data. Maps will typically be 100x100 pixels at about 500 meter resolution. At a range of 20,000 km, using the camera above, the pixel size would be 1 km, with the frame covering 1000 km. If Mars were tiled with 1500 landmarks, there would be about ten such landmarks in the field of view. A range error of ten kilometers would cause an easily detectible 0.5 pixel shift in the relative locations of two landmarks 1000 pixels apart. Furthermore, if the camera were not nadir pointed, so that landmarks were at different depths, then similar cross track errors could be detected. By using many images in the analysis, the spacecraft location could be determined to better than a kilometer. This processing could easily be accomplished on board the spacecraft, without downloading the images at all.

Figure 5 shows projected images and illuminated reconstructions near the potential Mars 2003 landing site in Isidis. The area shown is 150 km square, somewhat larger than the maps proposed above. Both MOC [4] and Viking images were used in the reconstruction.




Figure 5. Projected MOC and Viking images (top row) and map reconstructions near potential Isidis landing site

Since the MGS navigation is very accurate, the only unknowns are the landmark locations and topography, and the offset between the MOC pointing and the spacecraft frame. The Viking data is included to allow for stereo-photoclinometric analysis. Ultimately, the solutions will result in refinements to the Viking Orbiter navigation.

The digital elevation maps discussed in section 3 can be re-imaged to show the landmark at any viewing angle and illumination, as shown in figure 6. In addition to characterizing the actual surface, these maps are also being used as starting points for the creation of artificial terrain, consistent with the actual surface, for use in landing and rover studies [5] .




Figure 6. Image reconstructed from Isidis DEM.

Acknowledgements

This work was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

References

- [1] W. M. Owen, private communication.
- [2] R. W. Gaskell, "Digital Identification of Cartographic Control Points," Photogrammetric Engineering and Remote Sensing, 54:(6), 723-727, Part 1, June 1988.
- [3] B. Hapke, "Bidirectional Reflectance Spectroscopy, I. Theory," Journal of Geophysical Research, 86: (84), 3039-3054, April, 1981.
- [4] M. C. Malin, K. S. Edgett, M. H. Carr, G. E. Danielson, M. E. Davies, W. K. Hartmann, A. P. Ingersoll, P. B. James, H. Masursky, A. S. McEwen, L. A. Soderblom, P. Thomas, J. Veverka, M. A. Caplinger, M. A. Ravine, T. A. Soulanille, and J. L. Warren, MOC wide angle images M0101566 and M0104553 (www.msss.com/mars_images, 1999).
- [5] R. Gaskell, L. E. Husman, J. B. Collier and R. L. Chen, "Synthetic Environments for Simulated Missions," IEEE paper 13.0406, IEEE Aerospace Conference, Mar. 2001, Big Sky, Montana.